arXiv:1703.01973v1 [stat.ML] 6 Mar 2017

Batched High-dimensional Bayesian Optimization via Structural Kernel Learning

Zi Wang*† Chengtao Li*† Stefanie Jegelka† Pushmeet Kohli‡

Abstract

Optimization of high-dimensional black-box functions is an extremely challenging problem. While Bayesian optimization has emerged as a popular approach for optimizing black-box functions, its applicability has been limited to low-dimensional problems due to its computational and statistical challenges arising from high-dimensional settings. In this paper, we propose to tackle these challenges by (1) assuming a latent additive structure in the function and inferring it properly for more efficient and effective BO, and (2) performing multiple evaluations in parallel to reduce the number of iterations required by the method. Our novel approach learns the latent structure with Gibbs sampling and constructs batched queries using determinantal point processes. Experimental validations on both synthetic and real-world functions demonstrate that the proposed method significantly outperforms existing state-of-the-art approaches.

1 Introduction

Optimization is one of the fundamental pillars of modern machine learning. Considering that most modern machine learning methods involve the solution of some optimization problem, it is not surprising that many recent breakthroughs in this area have been on the back of more effective techniques for optimization. A case in point is deep learning whose rise has been mirrored by the development of numerous techniques like batch normalization.

While modern algorithms have been shown to be very effective for convex optimization problems defined over continuous domains, the same cannot be stated for non-convex optimization, which has generally been dominated by stochastic techniques. During the last decade, Bayesian optimization has emerged as a popular approach for optimizing black-box functions. However, its applicability is limited to low-dimensional problems because of computational and statistical challenges that arise from optimization in high-dimensional settings.

In the past, these two problems have been addressed by assuming a simpler underlying structure of the black-box function. For instance, Djolonga et al. (2013) assume that the function being optimized has a low-dimensional effective subspace, and learn this subspace via low-rank matrix recovery. Similarly, Kandasamy et al. (2015) assume additive structure of the function where different constituent functions operate on disjoint low-dimensional subspaces. The subspace decomposition can be partially optimized by searching possible decompositions and choosing the one with the highest GP marginal likelihood (treating the decomposition as a hyper parameter of the GP). Fully optimizing the decomposition is, however, intractable. Li et al. (2016) extended (Kandasamy et al., 2015) to functions with a projected-additive structure, and approximate the projective matrix via projection pursuit with the assumption that the projected subspaces have the same and known dimensions. The above-mentioned approaches share the computationally challenging problem of learning the groups of decomposed subspaces without assuming the dimension of the subspace is known. Both Kandasamy et al. (2015) and subsequently Li et al. (2016) adapt the decomposition by maximizing the GP marginal likelihood every certain number of iterations. However, such

*Authors contributed equally to this work.
†Massachusetts Institute of Technology. {ziw,ctli,stefje}@mit.edu
‡Microsoft Research. pkohli@microsoft.com

maximization is computationally intractable due to the combinatorial nature of the partitions over feature space, which forces prior work to adopt randomized search heuristics.

In this paper, we present a new formulation of Bayesian optimization specialized for high dimensions. One of the key contributions of this work is a new formulation that interprets prior work on high-dimensional Bayesian optimization (HDBO) through the lens of structured kernels, and places a prior on the kernel structure. Thereby, our formulation enables simultaneous learning of the decomposition of the function domain.

Prior work on latent decomposition of the feature space considers the setting where exploration/evaluation is performed once at a time. This approach makes Bayesian optimization time-consuming for problems where a large number of function evaluations need to be made, which is the case for high-dimensional problems. To overcome this restriction, we extend our approach to a batched version that allows multiple function evaluations to be performed in parallel (Desautels et al., 2014; González et al., 2016; Kathuria et al., 2016). Our second contribution is a principled approach to select the batch of evaluations for structured kernel learning-based HDBO.

1.1 Other Related Work

In the past half of a century, a series of different acquisition functions was developed for sequential BO in relatively low dimensions (Kushner, 1964; Močkus, 1974; Srinivas et al., 2012; Hennig & Schuler, 2012; Hernández-Lobato et al., 2014; Kawaguchi et al., 2015; Wang et al., 2016a; Kawaguchi et al., 2016). More recent developments address high dimensional BO, by making assumptions on the latent structure of the function to be optimized, such as low-dimensional structure (Wang et al., 2016b; Djolonga et al., 2013), or additive structure of the function (Li et al., 2016; Kandasamy et al., 2015). Learning such structure is, however, very challenging (Duvenaud et al., 2013).

While the aforementioned methods are sequential in nature, the growth of computing power has allowed parallel observation of function values, triggering new settings where function values of a *batch* of points are observed in one time step (Contal et al., 2013; Desautels et al., 2014; González et al., 2016). For example, in UCB-PE (Contal et al., 2013), the algorithm greedily chooses the points with highest posterior variances and updates the variances for other points for the next selection. It exploits the property of Gaussian Processes that the variance of posterior distribution does not depend on function mean. B-UCB (Desautels et al., 2014) takes a similar approach. When selecting a batch, the algorithm greedily chooses points with the highest UCB score computed via the out-dated function mean but up-to-date function variances. However, these methods may be too greedy in their selection, resulting in points that lie far from an optimum. Recent work (Kathuria et al., 2016) aims to resolve this issue by sampling the batch via a diversity-promoting distribution for better randomized exploration.

2 Background

Let $f : \mathcal{X} \to \mathbb{R}$ be an unknown function we aim to optimize over a compact set $\mathcal{X} \subseteq \mathbb{R}^D$. Within as few function evaluations as possible, we want to find

$$f(x^*) = \max_{x \in \mathcal{X}} f(x).$$

Following (Kandasamy et al., 2015), we assume a latent decomposition of the feature dimensions $[D] = \{1, \ldots, D\}$ into disjoint subspaces, namely, $\bigcup_{m=1}^{M} A_m = [D]$ and $A_i \cap A_j = \emptyset$ for all $i \neq j$, $i, j \in [D]$. Further, f can be decomposed into the following additive form:

$$f(x) = \sum_{m \in [M]} f_m(x^{A_m}).$$

To make the problem tractable, we assume that each f_m is drawn independently from $\mathcal{GP}(0, k^{(m)})$ for all $m \in [M]$. The resulting f will also be a sample from a GP: $f \sim \mathcal{GP}(\mu, k)$, where the priors are

$$\mu(x) = \sum_{m \in [M]} \mu_m(x^{A_m}), \quad k(x, x') = \sum_{m \in [M]} k^{(m)}(x^{A_m}, x'^{A_m}).$$

Let $\mathcal{D}_n = \{(x_t, y_t)\}_{t=1}^n$ be the data we observed from function f where $y_t \sim \mathcal{N}(f(x_t), \sigma)$. The log data likelihood for $\mathcal{D}_n$ is

$$\log p(\mathcal{D}_n | \{k^{(m)}, A_m\}_{m\in[M]}) = -\frac{1}{2}\boldsymbol{y}^{\mathrm{T}}(\boldsymbol{K}_n + \sigma^2\boldsymbol{I})^{-1}\boldsymbol{y} - \frac{1}{2}\log|\boldsymbol{K}_n + \sigma^2\boldsymbol{I}| - \frac{n}{2}\log 2\pi \tag{2.1}$$

where $\boldsymbol{K}_n = \left[\sum_{m=1}^M k^{(m)}(x_i^{A_m}, x_j^{A_m})\right]_{i\le n, j\le n}$ is the gram matrix associated with $\mathcal{D}_n$, and $\boldsymbol{y} = [y_t]_{t\le n}$ is the concatenated observed function values. Conditioned on the observations $\mathcal{D}_n$, we can infer the posterior mean and covariance function of the function component $f^{(m)}$ to be

$$\mu_n^{(m)}(x^{A_m}) = \boldsymbol{k}_n^{(m)}(x^{A_m})^\top(\boldsymbol{K}_n + \sigma^2\boldsymbol{I})^{-1}\boldsymbol{y},$$
$$k_n^{(m)}(x^{A_m}, x'^{A_m}) = k^{(m)}(x^{A_m}, x'^{A_m}) - \boldsymbol{k}_n^{(m)}(x^{A_m})^\top(\boldsymbol{K}_n + \sigma^2\boldsymbol{I})^{-1}\boldsymbol{k}_n^{(m)}(x'^{A_m}),$$

where $\boldsymbol{k}_n^{(m)}(x^{A_m}) = [k^{(m)}(x_t^{A_m}, x^{A_m})]_{t\le n}$.

We use regret to evaluate BO algorithms both in the sequential case and the batch evaluation case. For sequential evaluations, let $\tilde{r}_t = \max_{x\in\mathcal{X}} f(x) - f(x_t)$ denote the immediate regret at iteration t. We are interested in both the averaged cumulative regret $R_T = \frac{1}{T}\sum_t \tilde{r}_t$, and the simple regret $r_T = \min_{t\le T}\tilde{r}_t$. For batch evaluations, we denote $\tilde{r}_t = \max_{x\in\mathcal{X}, b\in[B]} f(x) - f(x_{t,b})$ as the immediate regret obtained by the batch at iteration t, where B is the batch size. The averaged cumulative regret of the batch setting is computed by $R_T = \frac{1}{T}\sum_t \tilde{r}_t$, and the simple regret $r_T = \min_{t\le T}\tilde{r}_t$. Note that the averaged cumulative regret provides an upper bound on the simple regret.

3 Learning additive kernel structure

We take a Bayesian view on the task of learning the latent structure of the GP kernel. The decomposition of the input space $\mathcal{X}$ will be learned simultaneously with optimization, as more and more data is observed. Our generative model draws mixing proportions $\theta \sim \textsc{Dir}(\alpha)$. Each dimension j is assigned to one out of M groups via the decomposition assignment variable $z_j \sim \textsc{Multi}(\theta)$. The objective function is then $f(x) = \sum_{m=1}^M f_i(x^{A_m})$, where $A_m = \{m : z_j = m)\}$ is the set of support dimensions for function f_i, and each f_i is drawn from a GP. Finally, given an input x, we observe $y \sim \mathcal{N}(f(x), \sigma)$. Figure 1 illustrates the corresponding graphical model.

Given the observed data $\mathcal{D}_n = \{(x_t, y_t)\}_{t=1}^n$, we obtain a posterior distribution over possible decompositions that we include in the BO process as explained below. The posterior distribution of the decomposition assignment vector z and θ is

$$p(z, \theta \mid \mathcal{D}_n; \alpha) \propto p(\mathcal{D}_n \mid z)p(z \mid \theta)p(\theta; \alpha)$$

Marginalizing over θ yields the posterior distribution of the decomposition assignment

$$\begin{aligned} p(z \mid \mathcal{D}_n; \alpha) &\propto p(\mathcal{D}_n \mid z)\int p(z \mid \theta)p(\theta; \alpha)\,\mathrm{d}\theta \\ &\propto p(\mathcal{D}_n \mid z)\frac{\Gamma(\sum_m \alpha_m)}{\Gamma(D + \sum_m \alpha_m)}\prod_m \frac{\Gamma(|A_m| + \alpha_m)}{\Gamma(\alpha_m)} \end{aligned}$$

where $p(\mathcal{D}_n|z)$ is the data likelihood (2.1) for the additive GP given a fixed structure defined by z. We learn the posterior distribution for the decomposition variable z via Gibbs sampling, choose the decomposition z among the samples that achieves the highest data likelihood, and then proceed with BO. The Gibbs sampler repeatedly draws coordinate assignments z_j according to

$$p(z_j = m \mid z_{\neg j}, \mathcal{D}_n;\ \alpha) \propto p(\mathcal{D}_n \mid z)p(z_j \mid z_{\neg j}) \propto p(\mathcal{D}_n \mid z)(|A_m| + \alpha_m) \propto e^{\phi_m},$$

where

$$\phi_m = -\frac{1}{2}\boldsymbol{y}^\top(\boldsymbol{K}_n^{(z_j=m)} + \sigma^2\boldsymbol{I})^{-1}\boldsymbol{y} - \frac{1}{2}\log\det(\boldsymbol{K}_n^{(z_j=m)} + \sigma^2\boldsymbol{I}) + \log(|A_m| + \alpha_m)$$

Figure 1: Graphical model for the structured Gaussian process; η is the hyperparameter of the GP kernel; z controls the decomposition for the input space.

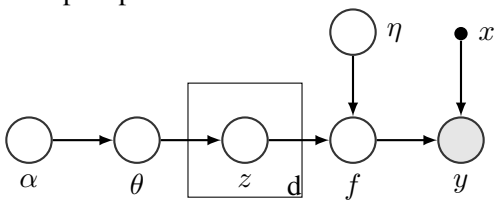

and $\boldsymbol{K}_n^{(z_j=m)}$ is the gram matrix associated with the observations $\mathcal{D}_n$ by setting $z_j = m$. We can use the Gumbel trick to efficiently sample from this categorical distribution. Namely, we sample a vector of i.i.d standard Gumbel variables ω_i of length M, and then choose the sampled decomposition assignment $z_j = \arg\max_{i \leq M} \phi_i + \omega_i$.

With a Dirichlet process, we could make the model nonparametric and the number M of possible groups in the decomposition infinite. Given that we have a fixed number of input dimension D, we set $M = D$ in practice.

4 Diverse Batch Sampling

In many real-world applications where function evaluations translate into time-intensive experiments, the typical sequential exploration strategy — observe one function value, update the model, then select the next observation via an acquisition function — is undesirable. *Batched Bayesian Optimization (BBO)* (Azimi et al., 2010; Contal et al., 2013; Kathuria et al., 2016) instead selects a batch of B observations in to be made in parallel before the model is updated with all simultaneously.

In such a scenario, the acquisition function by itself is not efficient to be optimized, nor does it sufficiently account for exploration. While the former issue is addressed by assuming an additive structure in the function, resolving the latter requires encouragement of observations which are informative by themselves but are also non-redundant considering other observations in the same batch. Recent work (Contal et al., 2013; Kathuria et al., 2016) achieves this goal by selecting a point that maximizes the acquisition function, and by adding the remaining points via a diversity criterion. In high dimensions, however, this diverse selection becomes expensive. To illustrate this, consider the case where each dimension has a finite number of possible values[1]. Here, the cost of sampling evaluations via a *Determinantal Point Process (DPP)*, as proposed in (Kathuria et al., 2016), grows exponentially with the number of dimensions due to the exponential growth of function domain. The same obstacle arises when employing the approach from (Contal et al., 2013), where points are sampled by greedily choosing the one with maximum posterior variance. Thus, naïve applications of these approaches in our setting would result in intractable algorithms. Instead, we propose a general approach that explicitly takes advantage of the structured kernel to enable relevant and non-redundant high-dimensional batch selection.

We now describe our approach for a single decomposition sampled from the posterior; it extends to a distribution of decompositions by sampling a set of decompositions from the posterior and then sampling points for each decomposition individually. Concretely, given a decomposition z, we define a Determinantal Point Process (DPP) on each group of dimensions individually. Assuming n observations have already been made, a set S of points in the subspace $\mathbb{R}^{|A_m|}$ is sampled with probability proportional to $\det(\mathbf{K}_n^{(m)}(S))$, where $\mathbf{K}_n^{(m)}$ is the posterior covariance matrix of the m-th group after observing n evaluations and $\mathbf{K}(S)$ is the submatrix of $\mathbf{K}$ with rows and columns indexed by S. Assuming the group sizes are upper-bounded by some constant, sampling from each such DPP individually implies an exponential speedup compared to using the full kernel.

[1]While we use this discrete categorical domain to illustrate the batch setting, our proposed method is general and is applicable to continuous box-constrained domains.

Sampling vs. Greedy Maximization A DPP on $\mathbf{K}_n^{(m)}$ assigns more probability mass to subset S with large $\det(\mathbf{K}_n^{(m)}(S))$, and this large determinant is equivalent to the set S being "diverse", i.e., the items in the set S are dissimilar, or, their feature vectors are fairly orthogonal. While finding the set S that maximizes the determinant is NP-hard, an approximate maximizer can be found by a greedy strategy where, in each step, we greedily add the point that leads to the maximum increase in the determinant, given the already selected points. This strategy is referred to as Pure Exploration (PE) in (Contal et al., 2013). We use both sampling and the greedy strategy for selecting a diverse batch of evaluations.

Sample Combination Once we have selected a diverse set of $(B-1)$ points for each group of dimensions (i.e., subspace), it remains to combine those points into query points in the joint space. Assume that for each group $m \in [M] = \{1, \ldots, M\}$ we have samples $\mathcal{X}_m = \{x_i^{(m)}\}_{i \in [B-1]}$ where $x_i^{(m)} \in \mathbb{R}^{|A_m|}$. There are exponentially many — $(B-1)^M$ — ways to combine the subspace points. We use two different methods for combinations.

First, a simple combination is *random selection without replacement*: sample one $x_i^{(m)}$ from each $\mathcal{X}_m$ uniformly randomly without replacement, and combine the $\{x_i^{(m)}\}_{m \in [M]}$ to get one sample in $\mathbb{R}^d$. We repeat this procedure until we obtain all the $(B-1)$ points to be observed. This retains diversity across the batch of samples since the samples are diverse within each group of features.

Second, we define a *quality function* $\psi_t^{(m)}$ for each group $m \in [M]$ at time t, and greedily combine samples to obtain high quality values $\psi_t(x) = \sum_m \psi_t^{(m)}$: in each of $B-1$ iterations, in each group we select $x_*^{(m)} = \arg\max_{x^{(m)} \in \mathcal{X}_m} \psi_t^{(m)}(x^{(m)})$, combine the selected subspace points, and remove them from the pool of selectable points: $\mathcal{X}_m \leftarrow \mathcal{X}_m \backslash \{x_*^{(m)}\}$.

Example: Add-UCB-DPP-BBO Next, we demonstrate the batch sampling for both the acquisition and the quality functions being *Upper Confidence Bounds (UCB)* (Srinivas et al., 2012). For iteration t and group m, one may define upper and lower confidence bounds on the unknown function value $f_t^{(m)}(x)$ as

$$\begin{aligned} (f_t^{(m)})^+(x) &= \mu_{t-1}^{(m)}(x) + \beta_t^{1/2}\sigma_t^{(m)}(x); \\ (f_t^{(m)})^-(x) &= \mu_{t-1}^{(m)}(x) - \beta_t^{1/2}\sigma_t^{(m)}(x), \end{aligned} \tag{4.1}$$

respectively, where β_t are appropriate constants. Following the UCB criterion, we set both the acquisition function $\phi_t^{(m)}$ and quality function $\psi_t^{(m)}$ to be $(f_t^{(m)})^+$ for group m at time t.

To respect the acquisition function in the batch sampling, we follow (Contal et al., 2013; Kathuria et al., 2016) and restrict the diverse sampling to a relevance region. The relevance region $(\mathcal{R}_t^{(m)})^+$ for group m in iteration t is

$$(\mathcal{R}_t^{(m)})^+ = \left\{ x \in \mathcal{X}_m \mid \mu_{t-1}^{(m)} + 2\sqrt{\beta_{t+1}^{(m)}}\sigma_{t-1}^{(m)} \geq (y_t^{(m)})^\bullet \right\},$$

where $(y_t^{(m)})^\bullet = (f_t^{(m)})^-((x_t^{(m)})^\bullet)$ and $(x_t^{(m)})^\bullet = \arg\max_{x^{(m)} \in \mathcal{X}_m} (f_t^{(m)})^-(x^{(m)})$. We then use $(\mathcal{R}_t^{(m)})^+$ as the ground set to sample with PE/DPP. The full algorithm is shown in Alg. 1, where the choice of acquisition and quality functions can be replaced by any desired criterion besides UCB.

5 Empirical Results

In this section, we show the empirical results of our proposed solutions for HDBO. The first part of this section focuses on verifying the effectiveness of using our Gibbs sampling algorithm to learn the decomposition of the additive function, and the second part presents the results on using batch Bayesian optimization for high dimensional problems with the Gibbs sampler.

Algorithm 1 Add-UCB-DPP-BBO Variants

Input: $\mathcal{X}$, N_{init}, N_{cyc}, T, B, M
Observe function values of N_{init} points chosen randomly from $\mathcal{X}$
Get the initial decomposition of feature space via Gibbs sampling and get corresponding $\mathcal{X}_m$'s
for $t = 1$ to T **do**
  if $(t \mod N_{cyc} = 0)$ **then**
    Learn the decomposition via Gibbs sampling and get corresponding $\mathcal{X}_m$'s
  end if
  Choose x_0 by maximizing UCB (acquisition function) for each group and combine them
  for $m = 1$ to M **do**
    Compute $(\mathcal{R}_t^{(m)})^+$ and $\mathbf{K}_{(t-1)B+1}^{(m)}$
    Sample $\{x_i^{(m)}\}_{i\in[B-1]} \subseteq \mathcal{X}_m$ via PE or DPP with kernel $\mathbf{K}_{(t-1)B+1}^{(m)}$
  end for
  Combine $\{x_i^{(m)}\}_{i\in[B-1],m\in[M]}$ either randomly or by maximizing UCB (quality function) without replacement to get $\{x_i\}_{i\in[B-1]}$
  Observe (noisy) function values for $\{x_i\}$ for $i \in \{0, \ldots, B-1\}$.
end for

Table 1: Empirical posterior of any two dimensions correctly being grouped together by Gibbs sampling.

d_x \ N	50	150	250	350	450
5	0.81 ± 0.28	0.91 ± 0.19	1.00 ± 0.03	0.97 ± 0.08	1.00 ± 0.00
10	0.21 ± 0.13	0.54 ± 0.25	0.68 ± 0.25	0.81 ± 0.27	0.93 ± 0.15
20	0.06 ± 0.06	0.11 ± 0.08	0.20 ± 0.12	0.43 ± 0.17	0.71 ± 0.22
50	0.02 ± 0.03	0.02 ± 0.02	0.03 ± 0.03	0.04 ± 0.03	0.06 ± 0.04
100	0.01 ± 0.01	0.01 ± 0.01	0.01 ± 0.01	0.01 ± 0.01	0.02 ± 0.02

5.1 Effectiveness of Decomposition Learning

We first probe the effectiveness of using the Gibbs sampling method described in Section 3 to learn the decomposition of the input space.

Recovering Decompositions First, we test on input dimension $D = 2, 5, 10, 20, 50, 100$ and functions sampled from a known additive Gaussian process prior with zero-mean and isotropic Gaussian kernel with bandwidth $= 0.1$ and scale $= 5$ for each function component. We randomly sample the decomposition settings which satisfy the constraint that there are at least two groups in the decomposition and at most 3 dimensions in each group.

We set the burn-in period to be 50 iterations, and the total number of iterations for Gibbs sampling to be 100. In Table 1 and Table 2, we show two quantities that are closely related to the learned empirical posterior of the decompositions with different numbers of randomly sampled observed data points (N). The first quantity, reported in Table 1, is the probability of two dimensions being correctly grouped together by Gibbs sampling in each iteration of Gibbs sampling after the burn-in period, namely, $\frac{\sum_{i<j\leq D} \mathbb{1}_{z_i^g \equiv z_j^g \wedge z_i \equiv z_j}}{\sum_{i<j\leq D} \mathbb{1}_{z_i \equiv z_j}}$. The second quantity, reported in Table 2, is the probability of two dimensions being correctly separated by Gibbs sampling in each iteration of Gibbs sampling after the burn-in period, namely, $\frac{\sum_{i<j\leq D} \mathbb{1}_{z_i^g \neq z_j^g \wedge z_i \neq z_j}}{\sum_{i<j\leq D} \mathbb{1}_{z_i \neq z_j}}$. The results show that the more observed data we have, the more accurate the learned decompositions are. It also verifies that the Gibbs sampling procedure can converge to the ground truth decomposition with enough data for relatively small number of dimensions. The higher the dimension, the more data we need to recover the true decomposition.

Table 2: Empirical posterior of any two dimensions correctly being separated by Gibbs sampling.

d_x \ N	50	150	250	350	450
2	0.30 ± 0.46	0.30 ± 0.46	0.90 ± 0.30	0.90 ± 0.30	1.00 ± 0.00
5	0.87 ± 0.17	0.80 ± 0.27	0.60 ± 0.32	0.55 ± 0.29	0.50 ± 0.34
10	0.88 ± 0.05	0.89 ± 0.06	0.89 ± 0.07	0.91 ± 0.08	0.94 ± 0.07
20	0.94 ± 0.02	0.94 ± 0.02	0.94 ± 0.02	0.95 ± 0.02	0.97 ± 0.02
50	0.98 ± 0.00	0.98 ± 0.00	0.98 ± 0.01	0.98 ± 0.00	0.98 ± 0.01
100	0.99 ± 0.00	0.99 ± 0.00	0.99 ± 0.00	0.99 ± 0.00	0.99 ± 0.00

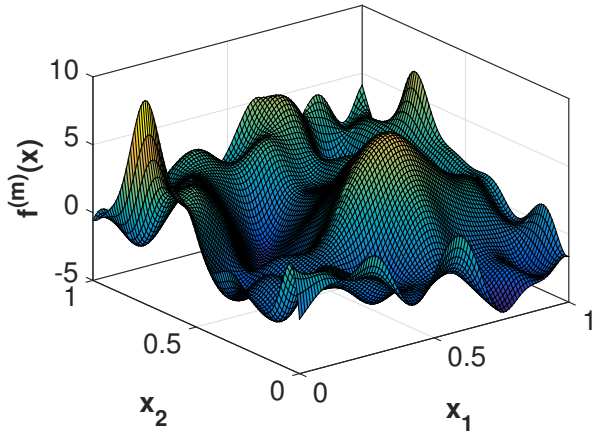

Figure 2: An example of a 2 dimensional function component of the synthetic function.

Effectiveness of Learning Decompositions for Bayesian Optimization To verify the effectiveness of the learned decomposition for Bayesian optimization, we tested on 2, 5, 10, 20, 30 and 50 dimensional functions sampled from a zero-mean add-GP with randomly sampled decomposition settings (at least two groups, at most 3 dimensions in each group) and isotropic Gaussian kernel with bandwidth $= 0.1$ and scale $= 5$. Each experiment was repeated for 50 times. An example of a 2 dimensional function component is shown in Fig. 2. Because of the numerous local maxima, it is very challenging to achieve the global optimum even for 2 dimensions, let alone maximizing an additive sum of them, only by observing their sum. For add-GP-UCB, we used $\beta_t^{(m)} = |A_m| \log 2t$ for lower dimensions ($D = 2, 5, 10$). and $\beta_t^{(m)} = |A_m| \log 2t/5$ for higher dimensions ($D = 20, 30, 50$). We show part of the results on averaged cumulative regret and simple regret in Fig. 3 and the rest in the appendix, comparing add-GP-UCB with known additive structure (`Known`), no partitions (`NP`), fully partitioned with one dimension for each group (`FP`) and the following methods of learning the decomposition: Gibbs sampling (`Gibbs`), random sampling the same number of settings of decompositions sampled by `Gibbs` and select the one with the highest data likelihood (`PL-1`), random sampling 5 settings of decompositions and select the one with the highest data likelihood (`PL-2`). For the latter two learning methods, the original add-GP-UCB method Kandasamy et al. (2015) denotes them as partial learning. The learning of the decomposition is done every 50 iterations. Fig. 4 shows the improvement of learning decompositions with `Gibbs` over optimizing without partitions (NP). From the results on the synthetic functions, it is clear that `Gibbs` outperforms both of the partial learning methods and for higher dimensions, `Gibbs` even out performs `Known`. This is because `Gibbs` may explore more in the beginning when the decomposition is not fully learned, which also facilitates finding the optimal value of the function.

We then tested the decomposition learning algorithm on a real-world function, which returns the distance between a designated goal location and two objects being pushed by two robot hands, whose trajectory is determined by 14 parameters specifying the location, rotation, velocity, moving direction etc. This func-

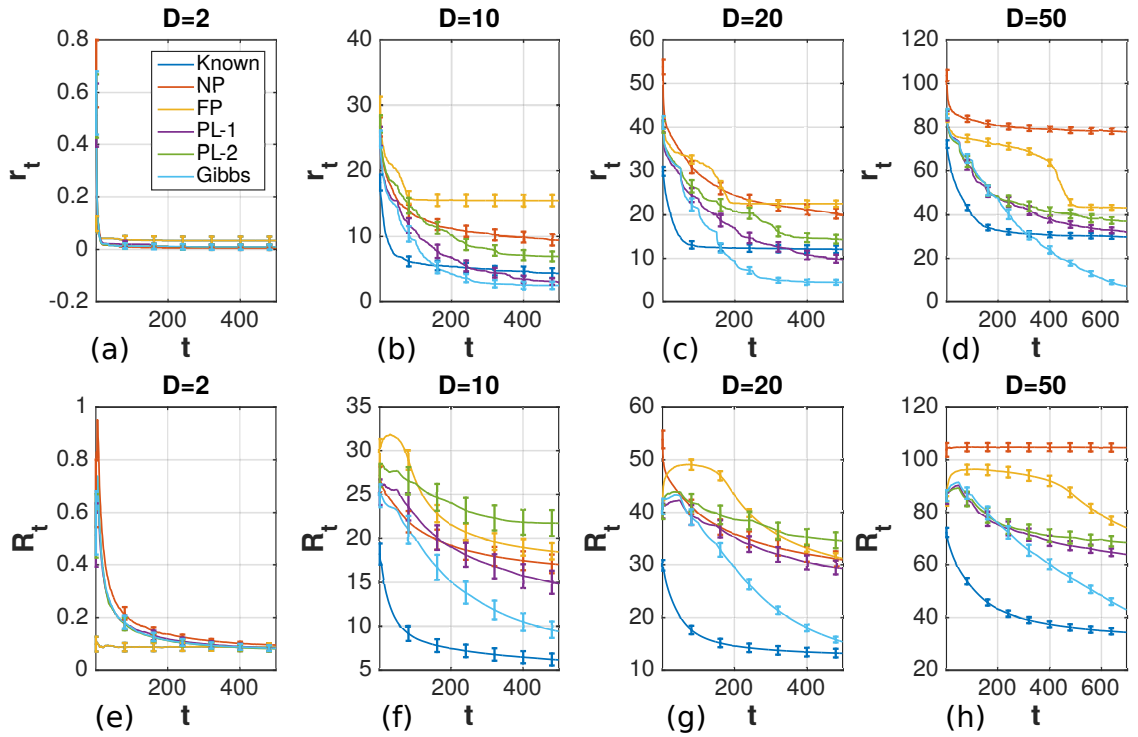

Figure 3: The simple regrets (r_t) and the averaged cumulative regrets (R_t) for setting input space decomposition with `Known`, `NP`, `FP`, `PL-1`, `PL-2`, and `Gibbs` on 2, 10, 20, 50 dimensional synthetic additive functions. `Gibbs` achieved comparable results to `Known`. Comparing `PL-1` and `PL-2` we can see that sampling more settings of decompositions did help to find a better decomposition. But a more principled way of learning the decomposition using `Gibbs` can achieve much better performance than `PL-1` and `PL-2`.

tion is inspired by Wang et al. (2017) and implemented with a physics engine, the Box2D simulator Catto (2011). We use add-GP-UCB with different ways of setting the additive structure to tune the parameters for the robot hand so as to push the object closer to the goal. The regrets are shown in Fig. 5. We observe that the performance of learning the decomposition with `Gibbs` dominates all existing alternatives including partial learning. Since the function we tested here is composed of the distance to two objects, there could be some underlying additive structure for this function in certain regions of the input space, e.g. when the two robots hands are relatively distant from each other so that one of the robot only impacts one of the object. Hence it is possible for `Gibbs` to learn a good underlying additive structure and perform effective BO with the structures it learned.

5.2 Diverse Batch Sampling

In this section, we focus on batch Bayesian optimization under high-dimensional setting. We demonstrate the effectiveness of Add-UCB-DPP-BBO shown in Sec. 4, and compare variants of this method. Concretely, we consider experimenting with the following baseline and four variants of Add-UCB-DPP-BBO:

- `Rand`: Each point for exploration is chosen uniformly randomly from $\mathcal{X}$
- `Batch-UCB-`$\star$: $\star \in \{$`PE`, `DPP`$\}$. All acquisition functions are UCB (Eq. 4.1). Exploration is done via PE or DPP with posterior covariance kernels for each group. Combination is done uniformly randomly.
- $\star$`-Fnc`: $\star \in \{$`Batch-UCB-PE`, `Batch-UCB-DPP`$\}$. Combination is done by maximizing quality functions, which are UCB's in our case.

Note that in all our experimental settings, direct applications of existing BBO methods are intractable due to the large exploration space, so we do not include those baselines.

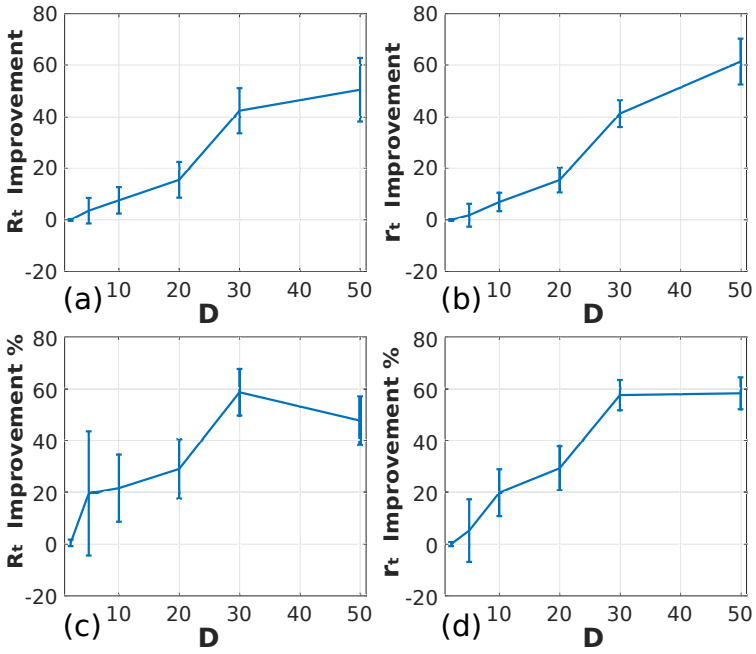

Figure 4: Improvement made by learning the decomposition with `Gibbs` over optimizing without partitions (NP). (a) averaged cumulative regret; (b) simple regret. (c) averaged cumulative regret normalized by function maximum; (d) simple regret normalized by function maximum. Using decompositions learned by `Gibbs` continues to outperform BO without `Gibbs`.

Effectiveness We tested on 2, 10, 20 and 50 dimensional functions sampled in the same way as in Sec. 5.1, where we assume the ground-truth decomposition of feature space is known. Since `Rand` performs the worst, we show relative averaged cumulative regret and simple regret of all methods against `Rand` in Fig. 6.

Results for absolute values of regrets are shown in the appendix. Each experiment was repeated for 20 times. For all experiments we set $\beta_t^m = |A_m| \log 2t$ and $B = 10$. All diverse batch sampling methods perform comparably well and far better than `Rand`, though there exist slight differences when comparing one against another. While in lower dimensional cases ($D \in \{2, 10\}$) `Batch-UCB-PE-Fnc` gives one of the best performances, in higher dimensional cases ($D \in \{20, 50\}$) `Batch-UCB-DPP-Fnc` performs better than (or comparable to) all other variants. When $D = 50$, the gap between `Batch-UCB-DPP-Fnc` and others becomes significant. We will see an even larger gap in later real-world experiments, showing that biasing the combination towards higher quality functions while retaining diversity across the batch of samples provides a better exploration-exploitation trade-off.

Fig. 7 shows the improvement of `Batch-UCB-DPP-Fnc` over `Rand` in terms of regrets. The improvement in general grows with d, indicating that the diverse batch sampling in higher dimensional settings will benefit more.

We also tested the diverse batch sampling algorithms for BBO on the `Walker` function which returns the walking speed of a three-link planar bipedal walker implemented in Matlab (Westervelt et al., 2007). We tune 25 parameters that may influence the walking speed, including 3 sets of 8 parameters for the ODE solver and 1 parameter specifying the initial velocity of the stance leg. We discretize each dimension into 40 points, resulting in a function domain of $|\mathcal{X}| = 40^{25}$, the size of which cannot be handled by direct application of existing batch sampling techniques. We learn the additive structure via Gibbs sampling and do batch sampling with $B = 5$. To enhance the algorithm efficiency, we limit the maximum size of each group of dimensions to be less than or equal to 2. The regrets for all methods are shown in Fig. 8. Again, all diverse batch sampling methods outperform `Rand` by a large gap. Moreover, `Batch-UCB-DPP-Fnc` gives significant better results compared to all its variants, supporting our claim that biasing the combination towards higher quality functions benefits more.

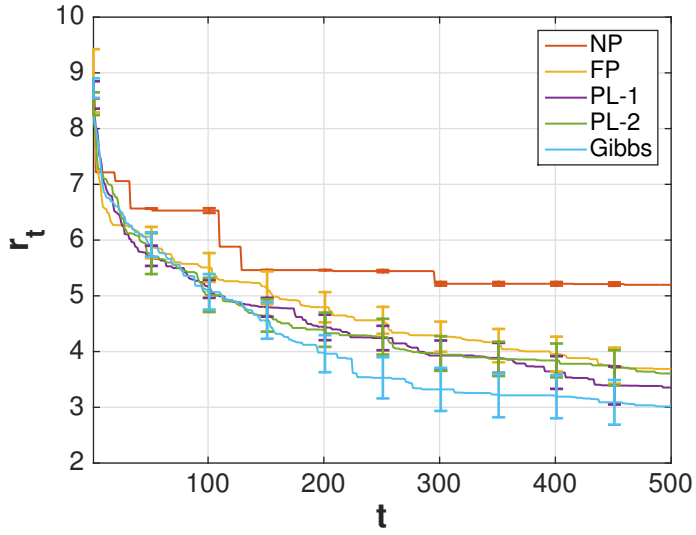

Figure 5: Simple regret of tuning the 14 parameters for a robot pushing task. Learning decompositions with `Gibbs` is more effective than partial learning (`PL-1`, `PL-2`), no partitions (`NP`), or fully partitioned (`FP`). Learning decompositions with `Gibbs` significantly helps BO for this tuning task.

Batch Sizes Finally, we test how the size of B affects the performances of proposed methods. We test the algorithms on a 14-dimensional `Robot` dataset with $B \in \{5, 10\}$. The regrets are shown in Fig. 5. In both settings `Batch-UCB-DPP-Fnc` achieves the best performance compared to all other variants. Meanwhile, increasing the batch size makes the gap to the other approaches more pronounced.

6 Conclusion

In this paper, we propose two novel solutions for high dimensional BO: inferring latent structure and performing batch optimization. The experimental results demonstrate that the proposed techniques are effective at optimizing high-dimensional black-box functions. Moreover, their relative performance against conventional methods grows higher with the dimensionality of the function. We believe that these results will lead to increased usage of Bayesian optimization for challenging optimization problems encountered in machine learning which typically involve a large number of parameters.

References

Azimi, Javad, Fern, Alan, and Fern, Xiaoli Z. Batch bayesian optimization via simulation matching. In *Advances in Neural Information Processing Systems (NIPS)*, 2010.

Catto, Erin. Box2D, a 2D physics engine for games. `http://box2d.org`, 2011.

Contal, Emile, Buffoni, David, Robicquet, Alexandre, and Vayatis, Nicolas. Parallel gaussian process optimization with upper confidence bound and pure exploration. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 225–240. Springer, 2013.

Desautels, Thomas, Krause, Andreas, and Burdick, Joel W. Parallelizing exploration-exploitation tradeoffs in Gaussian process bandit optimization. *Journal of Machine Learning Research*, 2014.

Djolonga, Josip, Krause, Andreas, and Cevher, Volkan. High-dimensional Gaussian process bandits. In *Advances in Neural Information Processing Systems (NIPS)*, 2013.

Duvenaud, David, Lloyd, James Robert, Grosse, Roger, Tenenbaum, Joshua B., and Ghahramani, Zoubin. Structure discovery in nonparametric regression through compositional kernel search. In *icml*, 2013.

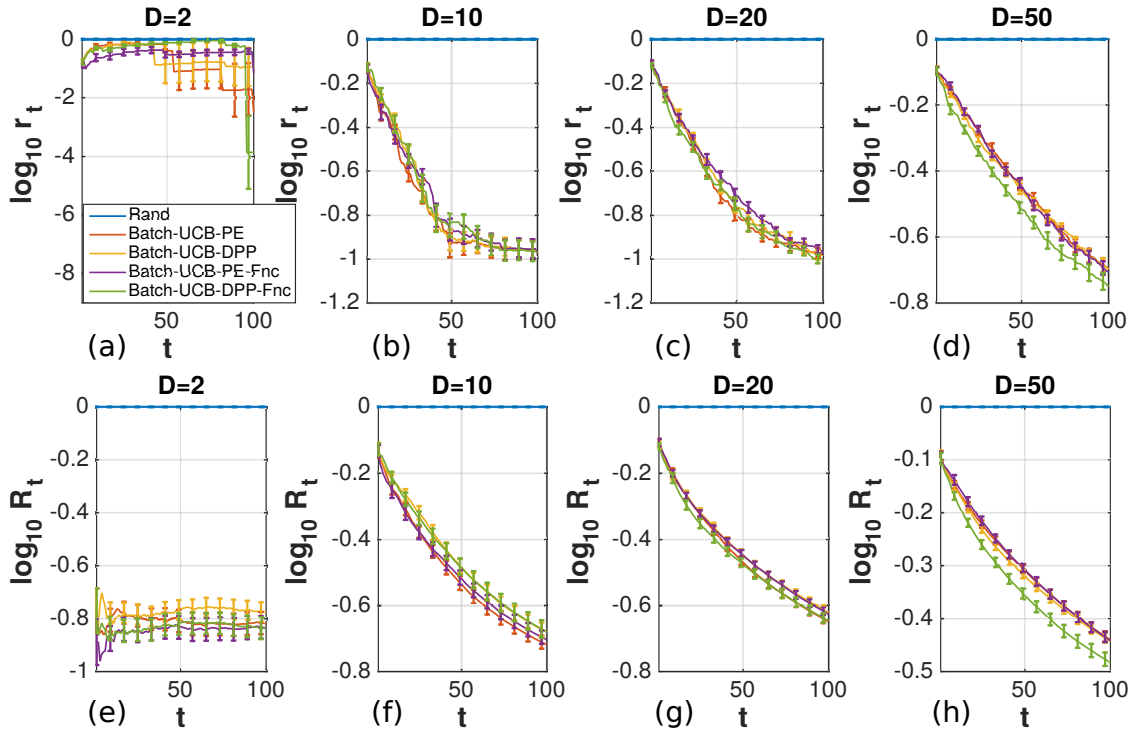

Figure 6: The scaled simple regrets (r_t) and the scaled averaged cumulative regrets (R_t) on synthetic functions with various dimensions when the ground truth decomposition is known. Four batch sampling methods (`Batch-UCB-PE`, `Batch-UCB-DPP`, `Batch-UCB-PE-Fnc` and `Batch-UCB-DPP-Fnc`) perform comparably well and outperform random sampling (`Rand`) by a large gap.

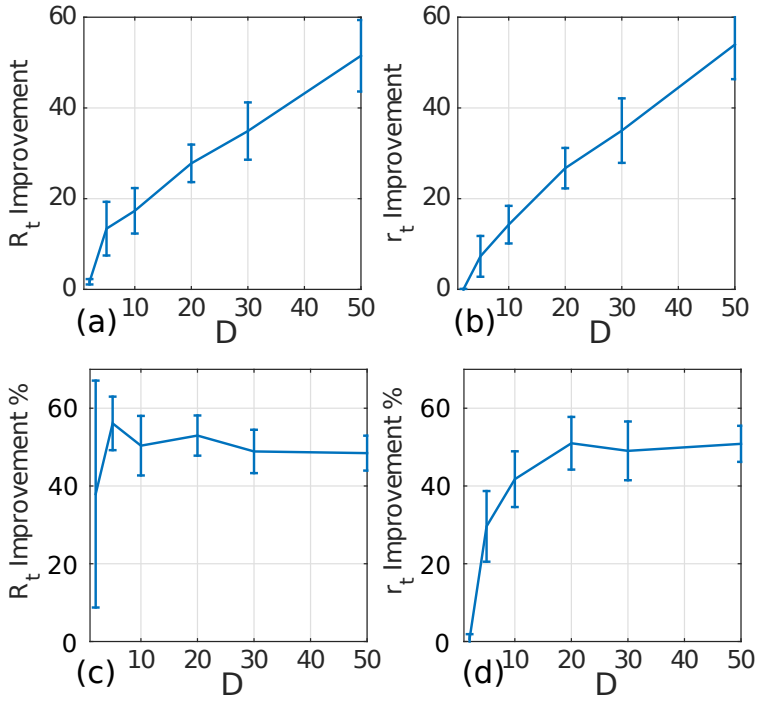

Figure 7: Improvement made efficient diverse batch sampling method `Batch-UCB-DPP-Fnc` over `Rand`. (a) averaged cumulative regret; (b) simple regret. (c) averaged cumulative regret normalized by function maximum; (d) simple regret normalized by function maximum. Using diverse batch sampling continues to outperform BBO with random exploration.

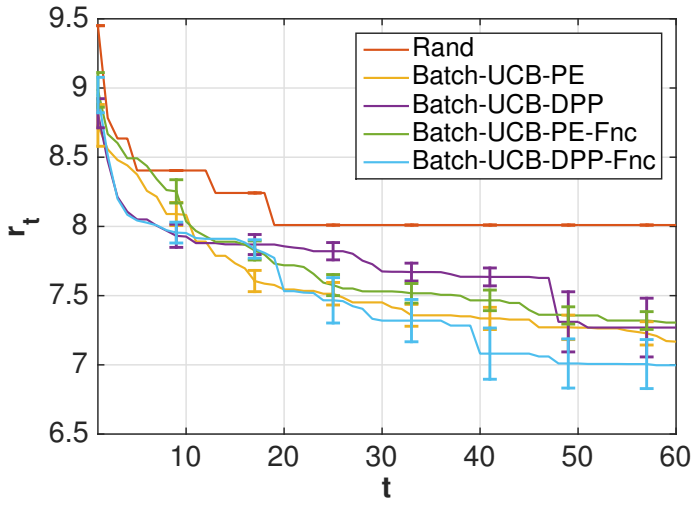

Figure 8: The simple regrets (r_t) of batch sampling methods on `Walker` data where $B = 5$. Four diverse batch sampling methods (`Batch-UCB-PE`, `Batch-UCB-DPP`, `Batch-UCB-PE-Fnc` and `Batch-UCB-DPP-Fnc`) outperform random sampling (`Rand`) by a large gap. `Batch-UCB-DPP-Fnc` performs the best among four diverse batch sampling methods.

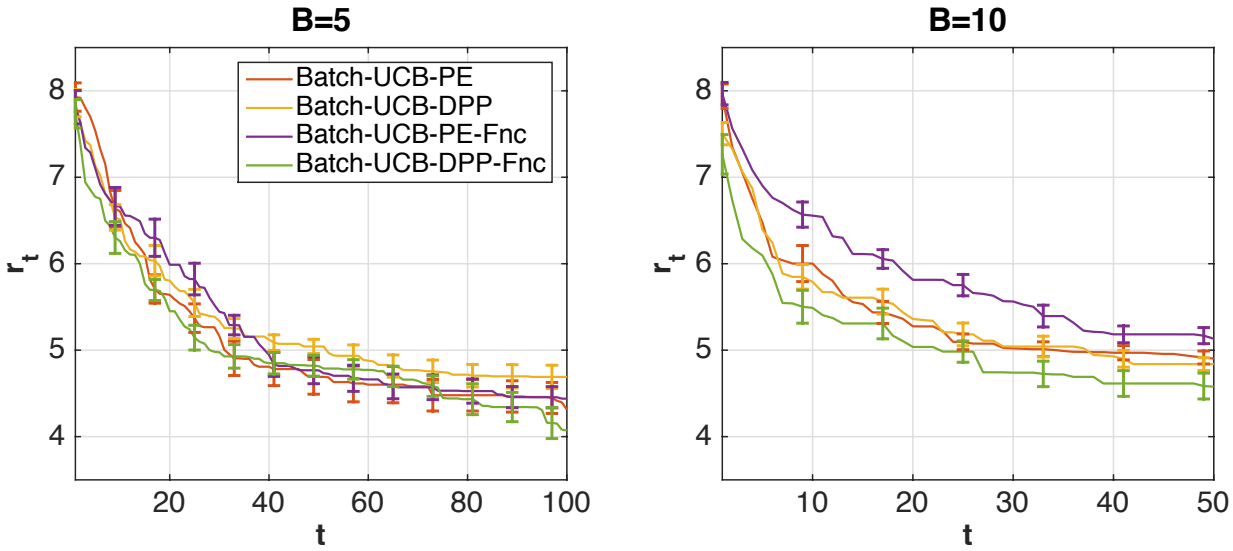

Figure 9: Simple regret when tuning the 14 parameters a robot pushing task with batch size 5 and 10. Learning decompositions with Gibbs sampling and diverse batch sampling are employed simultaneously to tackle the optimization problem. In general, `Batch-UCB-DPP-Fnc` performs best among four diverse batch sampling. The gap becomes more significant with batch size.

González, Javier, Dai, Zhenwen, Hennig, Philipp, and Lawrence, Neil D. Batch bayesian optimization via local penalization. *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2016.

Hennig, Philipp and Schuler, Christian J. Entropy search for information-efficient global optimization. *Journal of Machine Learning Research*, 13:1809–1837, 2012.

Hernández-Lobato, José Miguel, Hoffman, Matthew W, and Ghahramani, Zoubin. Predictive entropy search for efficient global optimization of black-box functions. In *Advances in Neural Information Processing Systems (NIPS)*, 2014.

Kandasamy, Kirthevasan, Schneider, Jeff, and Poczos, Barnabas. High dimensional Bayesian optimisation and bandits via additive models. In *International Conference on Machine Learning (ICML)*, 2015.

Kathuria, Tarun, Deshpande, Amit, and Kohli, Pushmeet. Batched Gaussian process bandit optimization via determinantal point processes. In *Advances in Neural Information Processing Systems (NIPS)*, 2016.

Kawaguchi, Kenji, Kaelbling, Leslie Pack, and Lozano-Pérez, Tomás. Bayesian optimization with exponential convergence. In *Advances in Neural Information Processing Systems*, pp. 2809–2817, 2015.

Kawaguchi, Kenji, Maruyama, Yu, and Zheng, Xiaoyu. Global continuous optimization with error bound and fast convergence. *Journal of Artificial Intelligence Research*, 56(1):153–195, 2016.

Kushner, Harold J. A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise. *Journal of Fluids Engineering*, 86(1):97–106, 1964.

Li, Chun-Liang, Kandasamy, Kirthevasan, Póczos, Barnabás, and Schneider, Jeff. High dimensional bayesian optimization via restricted projection pursuit models. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2016.

Močkus, J. On Bayesian methods for seeking the extremum. In *Optimization Techniques IFIP Technical Conference*, 1974.

Srinivas, Niranjan, Krause, Andreas, Kakade, Sham M, and Seeger, Matthias W. Information-theoretic regret bounds for gaussian process optimization in the bandit setting. *IEEE Transactions on Information Theory*, 2012.

Wang, Zi, Zhou, Bolei, and Jegelka, Stefanie. Optimization as estimation with Gaussian processes in bandit settings. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2016a.

Wang, Zi, Jegelka, Stefanie, Kaelbling, Leslie Pack, and Lozano-Pérez, Tomás. Focused model-learning and planning for non-gaussian continuous state-action systems. *International Conference on Robotics and Automation (ICRA)*, 2017.

Wang, Ziyu, Hutter, Frank, Zoghi, Masrour, Matheson, David, and de Feitas, Nando. Bayesian optimization in a billion dimensions via random embeddings. *Journal of Artificial Intelligence Research*, 55:361–387, 2016b.

Westervelt, Eric R, Grizzle, Jessy W, Chevallereau, Christine, Choi, Jun Ho, and Morris, Benjamin. *Feedback control of dynamic bipedal robot locomotion*, volume 28. CRC press, 2007.

A Additional experiments

In this section, we provide more details on our experiments.

A.1 Effectiveness of Decomposition Learning

BO for Synthetic Functions The full results of the simple and cumulative regrets for the synthetic functions comparing add-GP-UCB with known additive structure (`Known`), no partitions (`NP`), fully partitioned with one dimension for each group (`FP`) and the following methods of learning partition: Gibbs sampling (`Gibbs`), random sampling the same number of partitions sampled by `Gibbs` and select the one with the highest data likelihood (`PL-1`), random sampling 5 partitions and select the one with the highest data likelihood (`PL-2`) are shown in Fig. 11. The learning was done every 50 iterations, starting from the first iteration. For $D = 20, 30$, it is quite obvious that when a new partition is learned from the newly observed data (e.g. at iteration 100 and 150), the simple regret gets a boost.

BO for Real-world functions In addition to be 14 parameter robot pushing task, we tested on the walker function which returns the walking speed of a three-link planar bipedal walker implemented in Matlab Westervelt et al. (2007). We tune 25 parameters that may influence the walking speed, including 3 sets of 8 parameters for the ODE solver and 1 parameter specifying the initial velocity of the stance leg. To our knowledge, this function does not have an additive structure. The regrets of each decomposition learning methods are shown in Fig. 10. In addition to `Gibbs`, we test learning decomposition via constrained Gibbs sampling (`Gibbs-L`), where the maximum size of each group of dimensions does not exceed 2. Because the function does not have additive structure, `Gibbs` performed poorly since it groups together many dimensions of the input. As a result, its performance is similar to that of no partition (`NP`). However, `Gibbs-L` appears to learn a good decomposition with the group size limit, and manages to achieve a slightly lower regret than other methods. `Gibbs`, `PL-1`, `PL-2` and `FP` all performed relatively well in for this function, indicating that using the additive structure may benefit the BO procedure even if the function itself is not additive.

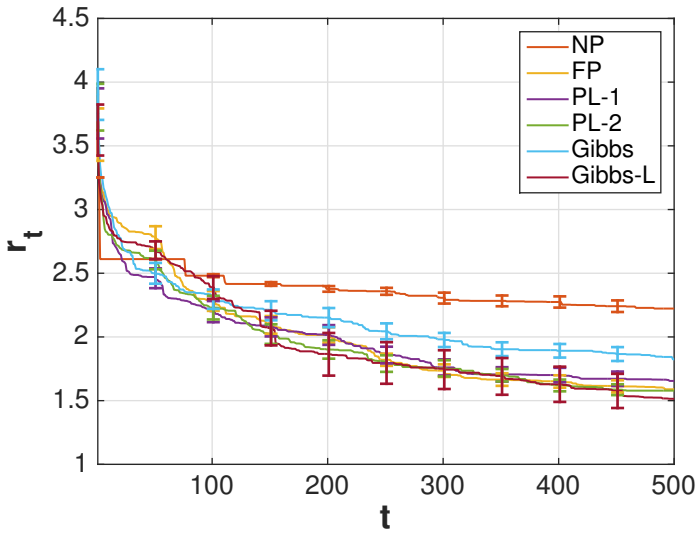

Figure 10: Simple regret of tuning the 25 parameters for optimizing the walking speed of a bipedal robot. We use the vanilla Gibbs sampling algorithm (`Gibbs`) and a Gibbs sampling algorithm with partition size limit set to be 2 (`Gibbs-L`) to compare with partial learning (`PL-1`, `PL-2`), no partitions (`NP`), and fully partitioned (`FP`). `Gibbs-L` performed slightly better than `PL-2` and `FP`. This function does not have an additive structure, and as a result, `Gibbs` does not perform well for this function because the sizes of the groups it learned tend to be large .

A.2 Diverse Batch Sampling

In Fig. 12, we show the full results of the simple and the cumulative regrets on the synthetic functions described in Section 5.2 of the paper.

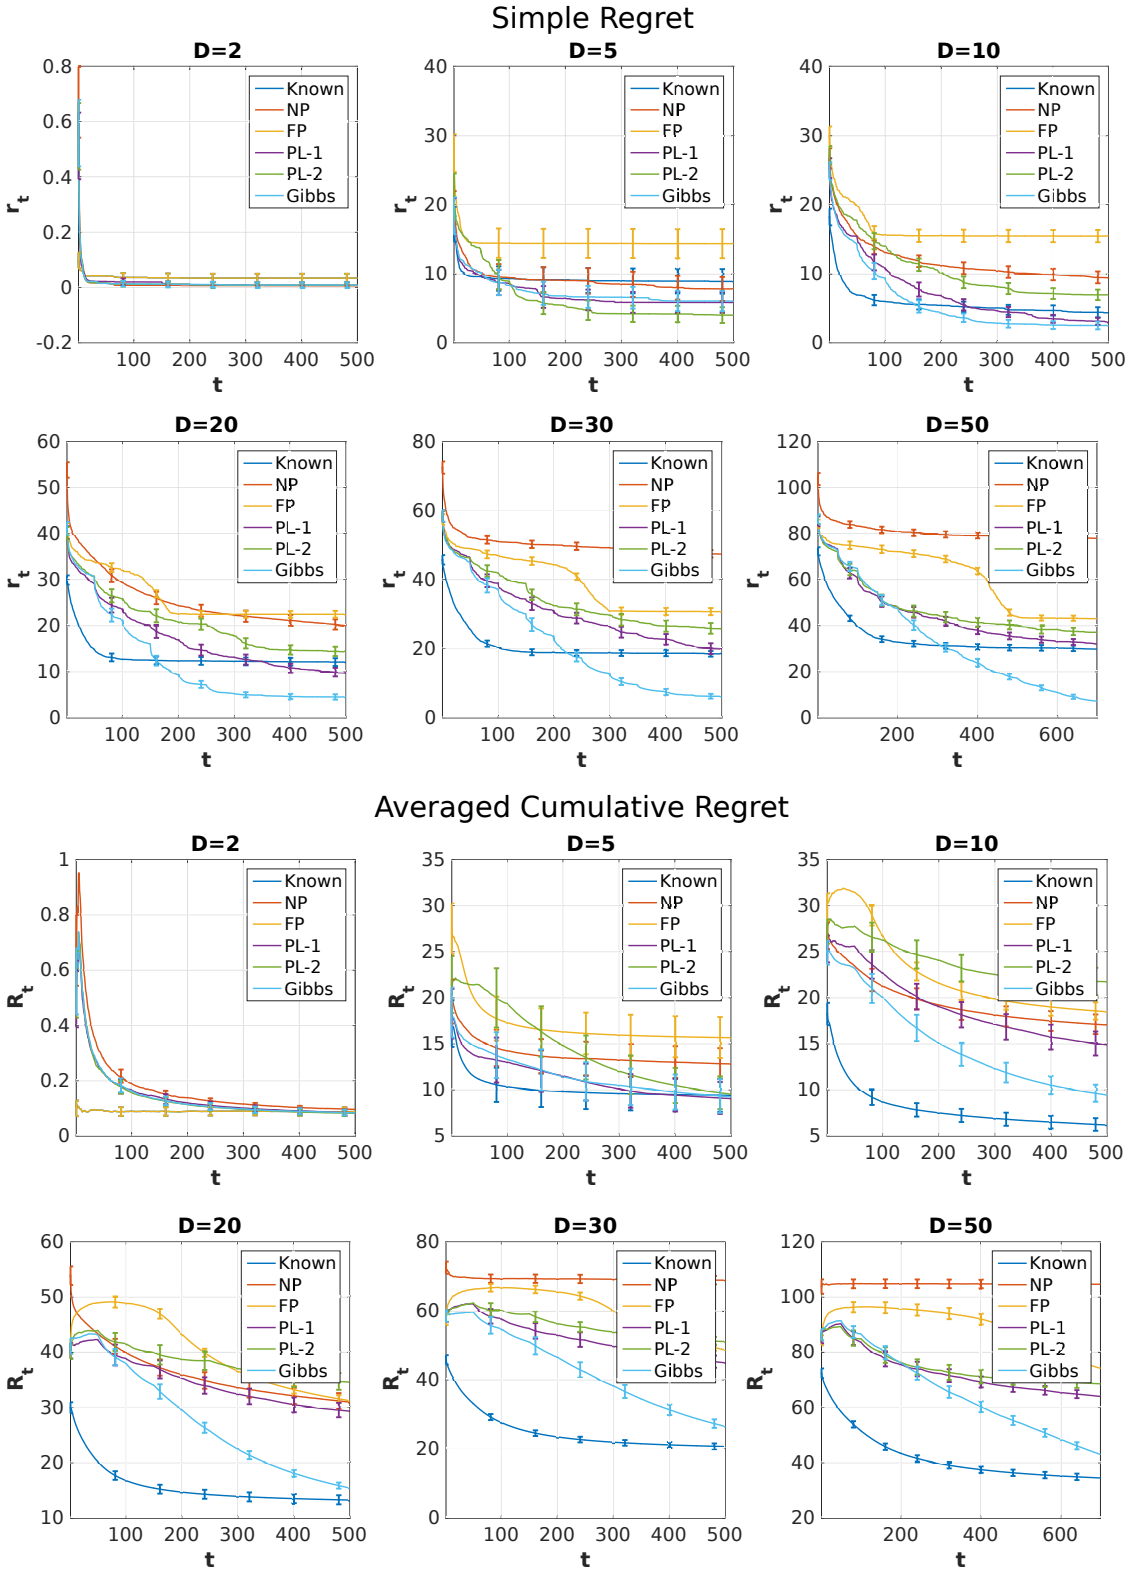

Figure 11: The simple regrets (r_t) and the averaged cumulative regrets (R_t) and for `Known` (ground truth partition is given), `Gibbs` (using Gibbs sampling to learn the partition), `PL-1` (randomly sample the same number of partitions sampled by `Gibbs` and select the one with highest data likelihood), `PL-2` (randomly sample 5 partitions and select the one with highest data likelihood), `FP` (fully partitioned, each group with one dimension) and `NP` (no partition) on 10, 20, 50 dimensional functions. `Gibbs` achieved comparable results to `Known`. Comparing `PL-1` and `PL-2` we can see that sampling more partitions did help to find a better partition. But a more principled way of learning partition using `Gibbs` can achieve much better performance than `PL-1` and `PL-2`.

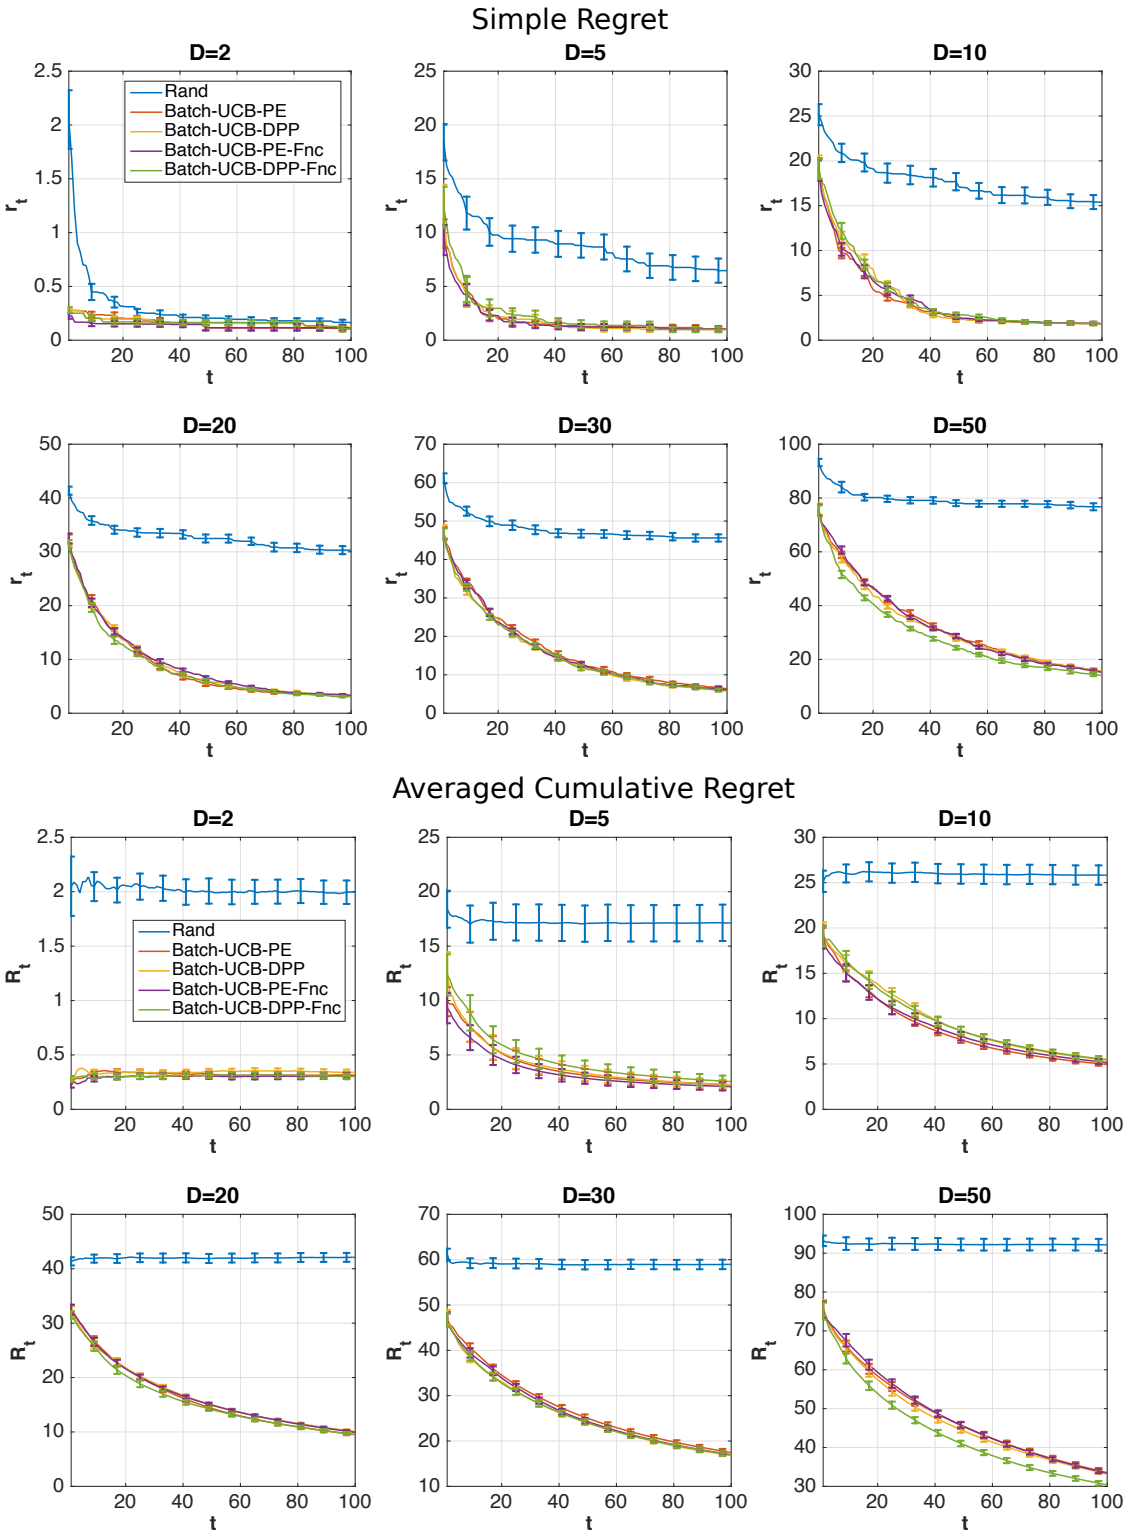

Figure 12: The simple regrets (r_t) and the averaged cumulative regrets (R_t) on synthetic functions with various dimensions when the ground truth partition is known. Four batch sampling methods (`Batch-UCB-PE`, `Batch-UCB-DPP`, `Batch-UCB-PE-Fnc` and `Batch-UCB-DPP-Fnc`) perform comparably well and outperform random sampling (`Rand`) by a large gap.